

## PRESERVING CACHE DATA AGAINST CLUSTER REBOOT

### ABSTRACT

A dual cluster storage server maintains track control blocks (TCBs) in a data structure to describe the data stored in cache in corresponding track images or segments. Following a cluster failure and reboot, the surviving cluster uses the TCBs to rebuild data structures such as a scatter table, which is a hash table that identifies a location of a track image, and a least recently used (LRU)/most recently used (MRU) list for the track images. This allows the cache data to be recovered. The TCBs describe whether the data in the track images is modified and valid, and describe forward and backward pointers for the data in the LRU/MRU lists. A separate non-volatile memory that is updated as the track images are updated is used to verify the integrity of the TCBs.